

AD-A122 599

AN INCENTIVE APPROACH TO ELICITING PROBABILITIES(U)
CALIFORNIA UNIV BERKELEY OPERATIONS RESEARCH CENTER
R D SHACHTER JUL 82 ORC-82-9 AFOSR-81-0122

1/1

UNCLASSIFIED

F/G 5/10

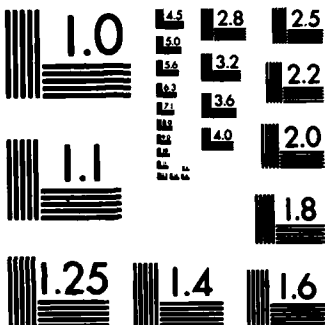
NL

END

FINED

+

DTIC



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

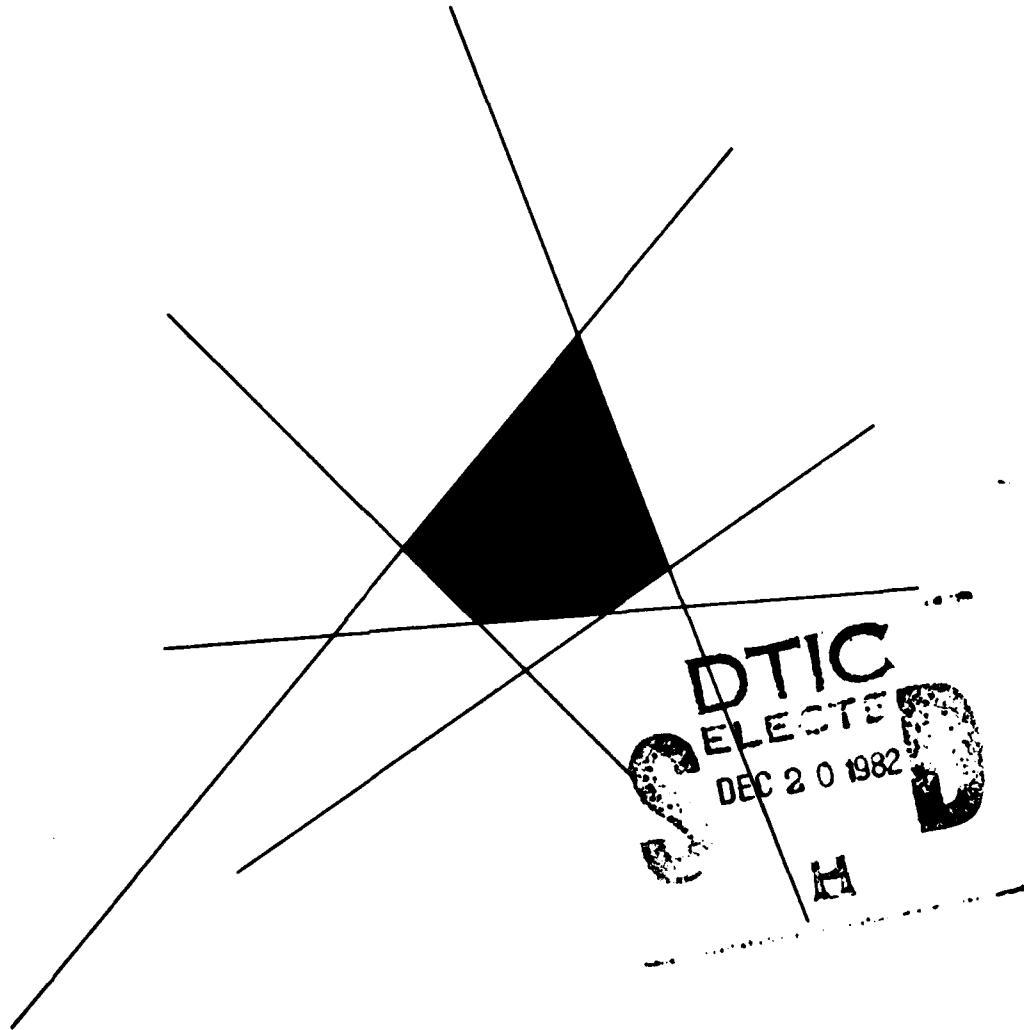
AD A 122599

ORC 82-9
JULY 1982

12

AN INCENTIVE APPROACH TO ELICITING PROBABILITIES

by
ROSS D. SHACTER



OPERATIONS
RESEARCH
CENTER

DTIC FILE COPY

UNIVERSITY OF CALIFORNIA • BERKELEY

12 82 006

AN INCENTIVE APPROACH TO ELICITING PROBABILITIES[†]

by

Ross D. Shachter
Department of Industrial Engineering
and Operations Research
University of California, Berkeley

JULY 1982

ORC 82-9

[†] Reproduction was supported by the Air Force Office of Scientific Research (AFSC), USAF, under Grant AFOSR-81-0122 with the University of California. Reproduction in whole or in part is permitted for any purpose of the United States Government.

To be published in the proceedings of the Workshop on Low Probability/High Consequence Risk Analysis by Plenum Press in 1983.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ORC 82-9	2. GOVT ACCESSION NO. AD-A122599	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) AN INCENTIVE APPROACH TO ELICITING PROBABILITIES		5. TYPE OF REPORT & PERIOD COVERED Research Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Ross D. Shachter		8. CONTRACT OR GRANT NUMBER(s) AFOSR-81-0122
9. PERFORMING ORGANIZATION NAME AND ADDRESS Operations Research Center University of California Berkeley, California 94720		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 2304/A5
11. CONTROLLING OFFICE NAME AND ADDRESS United States Air Force Air Force Office of Scientific Research Bolling Air Force Base, D.C. 20332		12. REPORT DATE July 1982
		13. NUMBER OF PAGES 14
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) Unclassified
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Probability Assessment Decision Analysis Risk Analysis Game Theory		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) (SEE ABSTRACT)		

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Abstract

A decision-maker (e.g., the Nuclear Regulatory Commission) seeks an expert's probabilities for uncertain quantities of interest (e.g., a seismologist's forecast of earthquakes), and wants the expert's reward to depend on the accuracy of the predictions. Assume that the expert compares compensation schemes on the basis of the expected utility of the dollar payoffs, and is willing to reveal his utility function for money. A reward is called 'proper' if the expert is never encouraged to state probabilities he does not truly believe. It is "strictly proper" if he is, in fact, encouraged to state his beliefs.

The reward procedure suggested in this paper uses the expert's stated probabilities and utility function to select from a set of possible payoffs. This procedure is always proper, but may not be strictly proper. If the preferred payoff is independent of the outcome whenever the decision-maker and expert agree on the probabilities, then they are said to be 'jointly risk-averse.' (For example, if the decision-maker agrees to play 'bookie' to a risk-averse expert, then they are jointly risk-averse.) In this case, the reward is shown to be strictly proper, as long as they don't disagree too much, so the expert can gain from researching the problem and carefully assessing his probabilities. In addition, the expert would prefer to make the bet more detailed, distinguishing between finer grain events, whenever such detail exposes new differences of opinion.

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



An Incentive Approach to Eliciting Probabilities

Ross D. Shachter

Department of Industrial Engineering and Operations Research
University of California, Berkeley

1. Introduction

A decision-maker wants to purchase an expert's probabilities for some uncertain events. As an example, consider the case of the Nuclear Regulatory Commission reviewing the license application for a nuclear power plant near a fault line in California. A critical factor in their analysis may be the probability of a major earthquake near the facility. They would like to contract with a seismologist to research the likelihood of such a quake.

The need for a procedure to reward the expert for his probabilities is illustrated by the following exchange. Professor Bruce Bolt at U. C. Berkeley has been a strong voice for earthquake awareness and preparation in California. He warns that [1] "... 50-50 odds can be viewed as a modest, perhaps even conservative appraisal of the likelihood that California will experience an earthquake of magnitude 7 or greater during the next ten years." Professor Richard Barlow of Berkeley has replied [2], "In many cases, though perhaps not in this instance, a probability of one-half is used when the forecaster has little knowledge concerning the event in question."

Prof. Barlow, trying to analyze the risks to the power plant, finds Prof. Bolt's statement too vague and fuzzy. There is no doubt regarding Bolt's qualifications to predict and study seismic activity. It would be useful, therefore, if he could offer a more exact statement of the probability.

In this paper, a procedure is developed that would encourage an expert to think about and express his precise beliefs when stating his probabilities. From a decision-theoretic viewpoint, a probability is defined in the context of comparing alternatives in an uncertain environment [3, 4]. When the expert is exposed to risk under uncertainty, he has a real need to think in terms of probabilities.

This suggests a payment to the expert for his probabilities that depends on what he states and what actually occurs. A "proper scoring rule" is a penalty function which the expert can minimize in expectation by assessing the probabilities he truly believes. A scoring rule is "strictly proper" if that minimum is unique. A probability can thus be defined with respect to a scoring rule [5, 6]. Necessary and sufficient conditions for a scoring rule to be strictly proper were given by Savage [7] and formalized by Haim [8].

The expert is assumed to have a utility function for money and limited assets, and to prefer that reward which maximizes his expected utility on this venture only. This will not coincide with minimizing his expected score unless the utility function is linear and the expert is quite wealthy, particularly when the amount of money involved is substantial. Therefore, if we desire a large enough incentive to encourage the expert to assess his probabilities, scoring rules will not work in general. It is necessary to incorporate the expert's utilities into the compensation scheme [9, 10, 11].

A reward will be called "proper" if the expert does not prefer to state probabilities he does not believe. A "strictly proper reward" encourages him to say what he believes. Rather than giving the expert a score to minimize, the procedure developed in this paper arrives at a reward function following negotiations between the decision-maker and the expert. Such a reward will always be proper, but it may not be strictly proper.

Consider a decision-maker and an expert who are not interested in betting on some events

when they agree on the probabilities for those events. Together they are called "jointly risk-averse." This is a generalization of the concept of risk-aversion for an individual [12, 13]. When they *don't* agree, they prefer to bet, and their disagreement is an economic good, from which the expert expects to gain. If they don't disagree too much, then the negotiated reward is strictly proper [14].

Suppose that the decision-maker would like to make the bet more detailed by splitting one of the events into two or more subevents. The expert prefers the new bet if it exposes new differences of opinion between them. This result can be used to obtain the expert's joint and marginal probability distributions for uncertain quantities.

2. Example of Procedure

In this section, we consider an example of a decision-maker purchasing an expert's probability for an uncertain event E , "a major earthquake will occur in California within the next ten years," and its complement \bar{E} . Suppose that the decision-maker expects to pay 20,000 dollars ten years from now for the expert's opinion, but is willing to pay up to 100,000 dollars for any outcome if necessary to obtain the expert's best assessment. The decision-maker thinks that the probability of an earthquake is $q = \frac{1}{3}$, and will agree to any payoff $x = (x_1, x_2) \in X$ where x_1 is the payment to the expert when an earthquake occurs and x_2 is the payment otherwise:

$$X = \{ x \in \mathbb{R}^2 : \frac{1}{3} x_1 + \frac{2}{3} x_2 \leq 20000; x_1 \leq 100000; x_2 \leq 100000 \}$$

The decision-maker may be thought of as playing the role of a bookie. The expert has 20,000 dollars to spend on bets. Each dollar bet on event E pays $\frac{1}{q} = 3$ dollars, should E occur; each dollar bet on event \bar{E} pays $\frac{1}{1-q} = \frac{3}{2}$ dollars if E does not occur.

The expert in this case is a seismologist, who we assume will follow the axioms of decision analysis. If he has utility function for money u , assets a_u , and believes that the probability of an earthquake is p , then he should select that payoff in X which maximizes his expected utility, given by

$$p u(a_u + x_1) + (1-p) u(a_u + x_2).$$

If we further assume that the expert would not like to risk any of his current holdings, then $a_u = 0$ and his expected utility for the preferred payoff is given by

$$\max \{ p u(x_1) + (1-p) u(x_2) : x \in X, x_1 \geq 0, x_2 \geq 0 \}.$$

The selling price or certainty equivalent of the preferred payoff is an amount such that the expert is indifferent between that amount for certain and the preferred payoff.

Consider a constantly risk-averse expert [13], with utility function

$$u(x) = -e^{-\alpha x}, x \geq 0.$$

In this case, if we apply Theorem 2 from Section 6,

$$x_1 = \max \{ 0, \min \{ 60000, 20000 + \frac{2}{3\alpha} \log \left(\frac{2p}{1-p} \right) \} \}$$

and

$$x_2 = \max \{ 0, \min \{ 30000, 20000 + \frac{1}{3\alpha} \log \left(\frac{1-p}{2p} \right) \} \}.$$

If he selects an interior point, then his expected utility is

$$p u(x_1) + (1-p) u(x_2) = u(20000) (3p)^{\frac{1}{3}} \left[\frac{3}{2} (1-p) \right]^{\frac{2}{3}} \geq u(20000)$$

with selling price

$$u^{-1}[p u(x_1) + (1-p) u(x_2)] = 20000 + \frac{1}{3\alpha} \log \left[(3p) \left(\frac{3}{2} (1-p) \right)^2 \right] \geq 20000.$$

The two inequalities are strict unless $p = q = \frac{1}{3}$. Thus, there are two cases. Either the decision-maker and expert agree on the probabilities in which case the expert receives 20,000 dollars for certain, or they disagree, and the expert receives more satisfaction by betting against the decision-maker and exploiting their difference of opinion.

In Table I, the expert's preferred payoffs are shown as a function of his probability p . The expert is assumed to have constant risk aversion $\alpha = 10^{-4}$. Note that there is a unique payoff for every probability p , $.0243 \leq p \leq .995$. The decision-maker, who believes that the probability is $\frac{1}{3}$, always expects to pay 20,000 dollars. The expert finds the payoff worth at least that much, and worth the most when he disagrees most with the decision-maker. Those are the times when his opinion may be considered most valuable to the decision-maker.

Table II shows the sensitivity of these results to the risk-aversion of the expert. When the expert is most risk-averse, his betting is timid, and his disagreement must be extreme before he will bet all 20,000 dollars. As a result, there is the widest range of probabilities for which his preferred payoff is unique. On the other hand, if the expert is least risk-averse, he may bet all he has even if his disagreement with the decision-maker is modest. Since the choice here is so sensitive to his probabilities and the stakes may be substantial, he should weigh his forecast carefully.

3. Reward and Payoff Functions

A decision-maker considers m mutually exclusive events E_1, \dots, E_m and their union, $F \equiv \bigcup_{i=1}^m E_i$, so that $\{E_1, \dots, E_m\}$ is a partition of F . The decision-maker would like the expert's reward to depend on both his stated probabilities and the observed event. The reward is then an uncertain venture, which entails the payment of x_i dollars when event E_i occurs (or receipt of $(-x_i)$ dollars if x_i is negative) and the m -vector payoff $x = (x_1, \dots, x_m)$ denotes the payment to the expert when event F occurs. When F does not occur, the decision-maker pays a side payment of c dollars.

Definition. A reward R is a function $R : P \rightarrow X$ where P is the set of probabilities

$$P = \{ p \in R^m : \sum_{i=1}^m p_i = 1; p_i \geq 0 \text{ for } i = 1, \dots, m \}$$

and $X \subset R^m$ is a compact set of possible payoffs to the expert. The expert states probabilities $p \in P$. If event F occurs, then exactly one event E_i occurs, and the decision-maker pays the expert $R_i(p)$ dollars. The real number c is a side payment if the decision-maker pays the expert c dollars when the event F does not occur.

Definition. The reward R is proper if the expert never prefers to state probabilities different from his true beliefs.

Definition. The reward R is strictly proper at probabilities p if the expert prefers to state his true beliefs when he believes p . A reward is strictly proper over a set if it is strictly proper at all points in the set.

Consider the constant reward $R(p) = (c, \dots, c)$. It is proper but not strictly proper.

Definition. A payoff function Y is a point-to-set mapping $Y : R \rightarrow X$ where $X \subset R^m$ is the set of possible payoffs. It will be assumed that $Y(c) \subseteq Y(d)$ if $c \leq d$.

Definition. R^c is an admissible reward if, should the expert believe probabilities p , there is no payoff in the set $Y(c)$ which he prefers to $R^c(p)$.

Claim. Admissible reward R^c is proper.

Table I. The expert's preferred payoff and his selling price for it as a function of his probability p . He has constant risk-aversion .0001 and the decision-maker expects to pay him \$20,000, but is willing to pay from \$0 to \$100,000.

p	1-p	Payoff		Selling Price
		E	not E	
0	1	0	30000	30000
.01	.99	0	30000	28253
.02	.98	0	30000	26767
.0243	.9757	3	29998	26190
.1	.9	9973	25014	22013
.2	.8	15379	22310	20487
.3	.7	18972	20514	20026
.3333	.6667	20000	20000	20000
.4	.6	21918	19041	20095
.5	.5	24621	17690	20566
.6	.4	27324	16338	21446
.7	.3	30270	14865	22850
.8	.2	33863	13069	25108
.9	.1	39269	10365	29337
.99	.01	55255	2372	44369
.995	.005	59910	45	48974
.996	.004	60000	0	50408
1	0	60000	0	60000

Table II. The expert's preferred payoff for different risk-aversions. The expert is constantly risk-averse and the decision-maker pays \$20,000 in expected value, and no less than zero dollars on any outcome.

Risk Aversion		.001	.0001	.00001	.000001
<hr/>					
Selling Price of a 50:50 chance at \$20,000 or \$0.		693	5662	9501	9950
Range of p for which expert does not bet all \$20,000.		1E-13 1-1E-26	.0243 .995	.270 .477	.3267 .3468
Probabilities .01, .99					
Payoff	E	17399	0	0	0
	not E	21301	30000	30000	30000
Selling Price		20905	28253	29651	29696
Probabilities .1, .9					
Payoff	E	18997	9973	0	0
	not E	20501	25014	30000	30000
Selling Price		20201	22013	26561	26959
Probabilities .3333, .6667					
Payoff	E	20000	20000	20000	20000
	not E	20000	20000	20000	20000
Selling Price		20000	20000	20000	20000
Probabilities .5, .5					
Payoff	E	20462	24621	60000	60000
	not E	19769	17690	0	0
Selling Price		20057	20566	25566	29550
Probabilities .99, .01					
Payoff	E	23526	55255	60000	60000
	not E	18237	2372	0	0
Selling Price		22437	44369	59181	59382

4. Assumptions of the Procedure

The procedure is based on the following assumptions about the behavior of the decision-maker and the expert relative to the mutually exclusive events E_1, \dots, E_m . It is hoped that in some contexts these assumptions will seem quite reasonable and the procedure may prove useful. There are situations where these assumptions are not so acceptable. In those cases, there may be no good way to reward an expert for his probabilities.

1. The decision-maker can set a point in time at which he will know which event has occurred.

This requires that the experiment has an unambiguous and indisputable outcome. The procedure seems most effective when that outcome will be known soon.

2. The expert is willing to bet up to a_u dollars, and the decision-maker a_v dollars on any specified outcome. The set of possible payoffs is

$$X = \{ x \in R^m : -a_u \leq x_i \leq a_v \text{ for } i = 1, \dots, m \}.$$

Both the decision-maker and expert should be willing to spend an amount independent of the outcome. While it is expected that one outcome may leave the decision-maker in a better asset position than another, the procedure is inherently a "zero-sum" game. This assumption eliminates the risk-sharing aspects possible with variable payoffs, in order to concentrate on the flow of information.

The expert's asset position, a_u , should not depend on the possible outcomes. As an example, consider our seismologist living near a major fault. He may agree completely with the decision-maker's probabilities and want to bet nonetheless as a form of earthquake insurance. It is theoretically possible to incorporate such conflict of interest in the analysis, but would be difficult to implement. It seems more reasonable to rely on the expert to be accurate and objective to establish and protect his reputation.

3. The expert believes probabilities $p = (p_1, \dots, p_m)$ and utility function $u : R \rightarrow R$ such that he prefers payoff $x \in X$ to payoff $y \in X$ if $\sum_{i=1}^m p_i u(a_u + x_i) > \sum_{i=1}^m p_i u(a_u + y_i)$. He is indifferent between them if $\sum_{i=1}^m p_i u(a_u + x_i) = \sum_{i=1}^m p_i u(a_u + y_i)$. The marginal utility u' is continuous and positive over $(0, a_u + a_v)$.

The expert's preferences between payoffs are characterized by his probabilities p and his utility function for money u . p_i is his subjective conditional probability for E_i given F , and it is unique. The utility function is not unique since any positive linear transformation $a u(x) + b$, $a > 0$, reveals the same preferences as u .

The assumption that there exists such a utility function of money is strong. The expert's satisfaction at the outcome is determined solely by the payment he will receive. If $x_i = x_j$, then the expert should not care whether it was event E_i or E_j which occurred. If the expert has a stake in the different outcomes or the decision-maker's decisions, then he may not be encouraged to reveal his true beliefs. While this assumption may be difficult for an expert to satisfy precisely, that may not be necessary. He is not in an adversary role with the decision-maker, and should be willing to cooperate in revealing his true probabilities, provided he does not suffer as a result.

The requirements on the marginal utility are for mathematical convenience. If u' is positive, then the expert is never saturated with money and always gets additional satisfaction from a larger payment.

4. The decision-maker offers payoff function

$$Y(c) = \{ x \in X : \sum_{i=1}^m q_i (a_v - x_i) \geq v(a_v - c) \}$$

where $-a_u < c < a_v$. He states non-zero probabilities $q = (q_1, \dots, q_m)$ and utility function $v : R \rightarrow R$, with marginal utility v' continuous and positive over $(0, a_u + a_v)$.

The decision-maker's q and v need not reflect his true beliefs. They are revealed at the start of the process before a side payment c is negotiated. If the expert finds them unacceptable, he may choose not to participate in the procedure. It is doubtful that the decision-maker's true beliefs could in fact be incorporated in a utility function independent of the outcome. If it is worth his while to consult the expert in the first place, then he most likely stands to gain or lose substantially from the events in question. Regardless, the expert, in selling his knowledge, should be insulated from such risks.

Ideally, q and v should be chosen to maximize the expected utility of the sample information. In practice that would be difficult to analyze. A reasonable q may be the decision-maker's true beliefs. If he thinks that the expert is "biased," then he may want to bias his q in the same direction. If the decision-maker chooses a linear function v then he is asking the expert to bet against a bookie. This is appealing in its simplicity and may be realistic if the decision-maker has considerably more assets than the expert, and is approximately risk-neutral over the range of possible payoffs. If, on the other hand, the decision-maker is conservative, and reluctant to pay much more than c dollars, he may want to state a risk-averse v . The effects of different risk-aversions on the admissible reward are shown in Table II.

5. The decision-maker and the expert agree to side payment c , $-a_u < c < a_v$, to be made at such time as it will be known which event has occurred.

This ensures that the decision-maker is willing to pay enough to obtain the expert's services.

Under these assumptions, the admissible reward $R^c(p)$ is a solution to

$$\begin{aligned} & \text{maximize } \sum_{i=1}^m p_i u(a_u + x_i) \\ & \text{subject to } \sum_{i=1}^m q_i v(a_v - x_i) \geq v(a_v - c) \\ & -a_u \leq x_i \leq a_v \text{ for } i = 1, \dots, m. \end{aligned}$$

This is precisely the "Pareto optimal" reward when two people bet on uncertain events in which they have no stake.

5. The Procedure

The procedure is now completely determined:

1. The decision-maker states events E_1, \dots, E_m , utility function v , assets a_v , and probabilities q . A point in time is set when it will be known which event has occurred so that the payoff can be made.
2. The expert states utility function u and assets a_u .
3. They agree to side payment c .
4. The expert states probabilities p .
5. After observing event E_i , the decision-maker pays the expert $R_i^c(p)$ dollars. If event F does not occur, then the decision-maker pays him c dollars.

Consider two examples where the decision-maker states a linear utility function and thus acts as a bookie. For simplicity, assume that they both have enough assets so that the admissible reward is in the interior of X . The admissible reward must satisfy the first order conditions for the mathematical program given in section 4. The general formula is given in Theorem 2 of section 6.

Suppose that the expert has constant risk aversion α , and utility function $u(x) = -e^{-\alpha x}$. Then

$$R_i^c(p) = c + \frac{1}{\alpha} \log \left[\frac{p_i}{q_i} \prod_{j=1}^m \left(\frac{p_j}{q_j} \right)^{-q_j} \right].$$

His expected utility is

$$\sum_{i=1}^m p_i u(R_i^c(p)) = u(c) \prod_{i=1}^m \left(\frac{p_i}{q_i} \right)^{q_i} \geq u(c)$$

(since $u(c) \leq 0$), and his selling price for the reward is

$$u^{-1} \left[\sum_{i=1}^m p_i u(R_i^c(p)) \right] = c + \frac{1}{\alpha} \log \left[\prod_{i=1}^m \left(\frac{q_i}{p_i} \right)^{q_i} \right] \geq c.$$

The two inequalities are strict unless $p_i = q_i$ for $i = 1, \dots, m$. Either they agree completely and the payment is c dollars no matter what happens, or the expert gains satisfaction by betting against the decision-maker. That gain comes from their difference of opinion.

Suppose instead that the expert has logarithmic utility, $u(x) = \log(x + s)$, $x \geq 0 > -s$. Then the admissible reward is

$$R_i^c(p) = \frac{p_i}{q_i} (c + s) - s$$

and his expected utility is

$$\sum_{i=1}^m p_i u(R_i^c(p)) = u(c) + \log \left[\prod_{i=1}^m \left(\frac{p_i}{q_i} \right)^{p_i} \right] \geq u(c).$$

That reward has selling price

$$u^{-1} \left[\sum_{i=1}^m p_i u(R_i^c(p)) \right] = (c + s) \prod_{i=1}^m \left(\frac{p_i}{q_i} \right)^{p_i} - s \geq c.$$

Again the two inequalities are strict unless $p_i = q_i$ for $i = 1, \dots, m$. They either agree completely and the reward is always c dollars, or the expert gains by expressing his disagreement.

In both of these examples the reward is constant when they agree and otherwise the expert gains satisfaction by betting. When v is linear, then any risk-averse expert should behave this way. In general, when this is true, the decision-maker and the expert will be called "jointly risk-averse."

6. Joint Risk-Aversion

Definition. The decision-maker and the expert are jointly risk-averse if the admissible reward when $p = q$ is unique, and $R^c(q) = (c, \dots, c)$.

Theorem 1.

The following statements are equivalent:

- The decision-maker and the expert are jointly risk-averse.
- $\frac{u'(a_u + c)}{v'(a_v - c)}$ is strictly decreasing in c .
- $\left[-\frac{u''(a_u + c)}{u'(a_u + c)} \right] + \left[-\frac{v''(a_v - c)}{v'(a_v - c)} \right] > 0$ for almost all c .

(Proofs for this theorem and the others appear in Shachter [14].)

Suppose that the decision-maker and the expert can agree on the probability that a given coin will land heads-up. They are jointly risk-averse if and only if they always prefer a constant payoff on the outcome of a flip. Otherwise, they would be interested in betting on the flip, even though they agree on the probabilities.

If the expert is risk-averse and the decision-maker states a linear or concave function v , they are jointly risk-averse. Even when the expert is risk-prone, they can still be jointly risk-averse if the decision-maker is sufficiently risk-averse.

Theorem 2.

Suppose that the decision-maker and the expert are jointly risk-averse. Then the following statements are true:

- Admissible reward $R^c(p)$ is a continuous function, for all c , $-a_u < c < a_v$:

$$R_i^c(p) = \max \{ -a_u, \min \{ a_v, b(\lambda \frac{q_i}{p_i}) \} \}$$

where b is defined by

$$b \left[\frac{u'(a_u + c)}{v'(a_v - c)} \right] = c$$

and λ is chosen such that

$$\sum_{i=1}^m q_i v(a_v - R_i^c(p)) = v(a_v - c).$$

- Given x , $-a_u < x_i < a_v$ for $i = 1, \dots, m$, then there is exactly one vector of probabilities p and side payment c such that $R^c(p) = x$.
- If $-a_u < R_i^c(p) < a_v$ for some i and c , then the expert prefers not to state probabilities r with $r_i \neq p_i$ when he believes p .
- If $-a_u < R_i^c(p) < a_v$ for $i = 1, \dots, m$, then R^c is strictly proper at p , and over a neighborhood of p .
- R^c is strictly proper at all p sufficiently close to q .

Consider an example in which the decision-maker and the expert are *not* jointly risk-averse. A simple case of this arises when they are both risk-neutral. If they are betting up to 1000 dollars each on a coin flip and the decision-maker considers the coin fair, then

$$\hat{R}(p) = \begin{cases} (-1000, 1000) & \text{if } p \leq .5 \\ (1000, -1000) & \text{if } p > .5 \end{cases}$$

is an admissible reward. So is

$$\tilde{R}(p) = \begin{cases} (-1000, 1000) & \text{if } p < .5 \\ (-500, 500) & \text{if } p = .5 \\ (1000, -1000) & \text{if } p > .5 \end{cases}$$

Note that the admissible reward is not unique. Also, the expert may always bet the 1000 dollars he has available, even when he agrees with the decision-maker. Finally, the admissible reward is not strictly proper, not even when the decision-maker and the expert agree on the probabilities.

7. More Details

Definition. Given $A \subseteq \{1, \dots, m\}$ and $E_A \equiv \bigcup_{i \in A} E_i$, the events $\{E_i : i \in A\}$ are said to be an **uninformative partition** of E_A if the decision-maker and expert agree on the conditional probability of E_i given E_A for all $i \in A$.

Consider $A = \{1, \dots, m\}$ so $E_A = F$, and $\{E_1, \dots, E_m\}$ is an uninformative partition of E_A if $p = q$, the decision-maker and the expert agree completely.

Theorem 3.

Suppose that the decision-maker and the expert are jointly risk-averse. Then the following statements are true:

- If $\{E_i : i \in A\}$ is an uninformative partition of E_A then $R_i^c(p) = R_j^c(p)$ for all i and $j \in A$.
- If $-a_u < R_i^c(p) < a_v$ for some i and c , and $R_i^c(p) = R_j^c(p)$ for all $j \in A$, then $\{E_i : i \in A\}$ is an uninformative partition of E_A .

Definition. Mutually exclusive, collectively exhaustive events D_1, \dots, D_n are called **reference events** for E_1, \dots, E_m if the decision-maker and the expert agree that E_i and D_j are independent (given F) for $i = 1, \dots, m$ and $j = 1, \dots, n$.

The decision-maker is not just interested in whether a major quake will occur, but would also like to know its location. Let event N be "the (first) major quake occurs in Northern California" and S be "the (first) major quake occurs in Southern California." Then instead of partition $\{E, \bar{E}\}$ of F , the decision-maker wants partition $\{EN, ES, \bar{E}\}$. The expert can never suffer by going to the more detailed bet, and will, indeed, prefer it when he disagrees with the decision-maker about the conditional probability of the location given that there is an earthquake.

In a similar fashion, the decision-maker needs information about the magnitude of a possible quake. Let event M_7 be "the (first) major quake has magnitude less than 8," and M_8 be "the (first) major quake has magnitude 8 or higher." The expert is willing to substitute $\{EM_7, EM_8, \bar{E}\}$ for $\{E, \bar{E}\}$, and prefers the switch if he disagrees with the decision-maker about the magnitude of possible earthquakes.

Now suppose that the decision-maker needs both location and magnitude information at the same time. Two methods of designing the multiple bet allow us to obtain either the marginal distributions or the joint probability distribution.

To assess the marginal distributions we need two reference events. Let H be "a flipped coin lands heads-up" and T be "a flipped coin lands tails-up." Clearly H and T are mutually exclusive and collectively exhaustive, and the outcome of the flip is independent of any earthquake. If the events $\{EHN, EHS, ETM_7, ETM_8, \bar{E}\}$ are used, then the decision-maker can obtain the expert's conditional distributions for magnitude and location at the same time. If they both agree on the probabilities of reference events H and T , then the expert's satisfaction is between that from the two separate bets. He does not reveal any of his opinions about their dependence. As more details about the events are added to the bet, the number of events grows slowly.

To assess the joint distributions we consider the events $\{ENM_7, ENM_8, ESM_7, ESM_8, \bar{E}\}$. The expert reveals his beliefs about the dependence of location and magnitude. The expert prefers this bet to all of the bets mentioned so far, if he disagrees with the decision-maker about that dependence. This is just a second application of Theorem 3. Note that the event space is the product of the magnitude and location events and the number of events can grow quickly as the bet gets more detailed.

8. Conclusions and Extensions

A procedure has been suggested to reward an expert in such a way that he is encouraged to reveal his true beliefs about uncertain quantities as probabilities. There is a negotiated fixed payment to him in exchange for his research time and effort and, in addition, he bets against the decision-maker. The expert prefers such a reward over the fixed payment alone, whenever he disagrees with the decision-maker's probabilities. The more they disagree, the more the expert expects to gain from the reward. The decision-maker, on the other hand, also prefers to uncover those disagreements, since he is interested in learning from the expert.

Once the decision-maker obtains the expert's probabilities, he can use Bayes' Law to revise his own opinion of the uncertain quantities. Let $P(E_i | F)$ be the decision-maker's prior probability for the event E_i given F . (Note that his stated probabilities q are not necessarily his true beliefs.) Let $P(p | E_i, q)$ be his likelihood function for the expert's stated probabilities p . The decision-maker's q is included because the expert may learn from it. Bayes' Law for the decision-maker's posterior probability is then

$$P(E_i | F, p, q) \propto P(p | E_i, q) P(E_i | F).$$

The decision-maker may wish to consult with more than one expert. This can be addressed by applying the procedure sequentially. In this way, his posterior distribution after each expert can be used as his prior distribution for the next. It can be used to determine whom, if anyone, he

should consult next and in the design of the next expert's reward. The likelihood for each expert changes as other experts are consulted. First, one would expect that experts on the same issues would be correlated, so that the decision-maker learns about the next expert from the previous ones [15]. Likewise, since the decision-maker incorporates the opinions of the experts already consulted, the next expert may have more confidence in the decision-maker's updated distribution.

The procedure is designed to assess probability for a finite number of disjoint events. That is extended to the case of overlapping events to obtain multiple marginal distributions. Another extension would be to an uncountable sample space. If both the decision-maker and expert express probability density functions which are nonzero, bounded and continuous over the sample space, then the admissible reward is of the same form as for finite events.

9. References

- [1] Bruce A. Bolt and Richard H. Jahns, California's Earthquake Hazard: A Reassessment, Public Affairs Report 20, 1-10, U.C. Berkeley (August 1979).
- [2] Richard E. Barlow, Assessment of Subjective Probability, Operations Research Center Report 81-23, U.C. Berkeley (1981).
- [3] Howard Raiffa, Decision Analysis: Introductory Lectures on Choices Under Uncertainty (Addison-Wesley, Reading, Mass., 1968).
- [4] P. C. Fishburn, Utility Theory, Management Science 14, 335-378 (1968).
- [5] Bruno de Finetti, Theory of Probability, Volume 1 (Wiley, New York, 1974).
- [6] D. V. Lindley, Scoring Rules and the Inevitability of Probability, International Statistical Review to appear (1982).
- [7] Leonard J. Savage, Elicitation of Personal Probabilities and Expectations, Journal of the American Statistical Association 66, 783-801 (1971).
- [8] Eduardo Haim, The Characterization of Strictly Proper Scoring Rules in Decision Making, Operations Research Center Report 81-22, U.C. Berkeley (1981).
- [9] Cedric A. B. Smith, Consistency in Statistical Inference and Decision, Journal of the Royal Statistical Society, Series B, 23, 1-37 (1961).
- [10] Robert L. Winkler, Scoring Rules and the Evaluation of Probability Assessors, Journal of the American Statistical Association 64, 1073-1078 (1969).
- [11] Robert L. Winkler and Allan H. Murphy, Nonlinear Utility and the Probability Score, Journal of Applied Meteorology 9, 143-148 (1970).
- [12] John W. Pratt, Risk Aversion in the Small and in the Large, Econometrica 32, 122-136 (1964).
- [13] Ralph L. Keeney and Howard Raiffa, Decisions with Multiple Objectives: Preferences and Value Tradeoffs (Wiley, New York, 1976).
- [14] Ross D. Shachter, The Economics of a Difference of Opinion: An Incentive Approach to Eliciting Probabilities, forthcoming PhD thesis, U.C. Berkeley (1982).
- [15] D. V. Lindley, A. Tversky, and R. V. Brown, On the Reconciliation of Probability Assessments, Journal of the Royal Statistical Society, Series A, 142, Part 2, 146-180 (1979).

END

FILMED

1-83